

REMARKS

Declaration

Applicant draws attention to the enclosed copy of a declaration signed by inventor Eitan Bachmat. This declaration was filed with the application, as indicated by the enclosed copy of the postcard stamped by the Office. Accordingly, no new declaration is required.

General remarks on the specification

The remarks made throughout the office action and the nature of the references selected suggest a fundamental misunderstanding of the claimed invention. Applicant offers the following analogy to assist the Office in better understanding the claimed invention.

Suppose one's home were heated with oil (i.e. the "incumbent" heating fuel) and one wished to consider whether or not to change to gas (i.e. the "competing" heating fuel). To decide whether or not to switch fuels, one might simulate what gas heat *would have cost* had the home been heated with gas instead of oil. The resulting "competing-fuel score" would then be compared with an "incumbent-fuel score," which would depend on the actual cost of oil used during the corresponding period.

Note that in the foregoing scenario, the home is *not* being heated by *both* gas and oil. It is *only* being heated by oil. Since the home is *not* actually being heated by gas, one cannot "evaluate" the cost of heating with gas; one must instead *simulate* that cost. This simulation might be carried out by, for example, inputting the desired indoor temperature, daily outdoor temperatures, winds, and home insulation parameters into a program provided to prospective customers by the local gas utility. This would result in a "competing-fuel score."

However, since the home is actually being heated by oil, there is no need to also simulate the cost of heating with oil. Instead, one could measure *actual data* from which one can derive an *actual cost* (i.e. "incumbent-fuel cost") to heat the home during any particular interval.

In the above example, one *simulates* the cost of heating with gas (i.e. the "competing" fuel) and one *evaluates* the cost of heating with oil (the "incumbent" fuel). The language used in

the claims was carefully selected to maintain this distinction. Hence, claim 1 refers to “*simulating* performance of a *competing* algorithm” but “*evaluating*...performance of an *incumbent* algorithm.”

Drawings

The Office objects to FIG. 2 because the source of the input data stream is not shown.

A drawing is required to show every feature of the invention specified in the claims (37 CFR 1.83(a)). In this case, the source of the input data stream is not specified in the claims. Hence, there is no reason to include it in FIG. 2.

Title

The Office's proposed title, “METHOD AND SYSTEM FOR EVALUATING PERFORMANCE OF ALGORITHMS EXECUTED BY A DATA-STORAGE SYSTEM,” suggests a fundamental misunderstanding of the nature of the invention. The claims do not recite evaluating the performances of various algorithms. They recite, in general, evaluating the performance of one algorithm, i.e. the “incumbent” algorithm, and simulating the performance of “competing” algorithms. The distinction between *evaluating* a performance and *simulating* a performance is important in understanding the nature of the invention.

Applicant submits that the present title “DYNAMIC DEMONSTRATION OF UNIMPLEMENTED ALGORITHMS” accurately and concisely states the nature of the invention as required by 37 CFR 1.72. The “unimplemented algorithms” referred to in the title are the competing algorithms recited in the claims. They are called “competing” because they compete with the “incumbent” algorithm. They are “unimplemented” because they are not actually executing. The competing-algorithm performance that results from simulating the competing algorithms amounts to a “demonstration” of what the performance of each competing algorithm would have been had it in fact been executing in place of the incumbent algorithm. The adjective “dynamic” modifies “demonstration” because the demonstration is being carried out in real-time on live data.

Objection to Abstract

The objection to the abstract and the section 112 rejection of claims 1, 10, 11, 13, 16, 21, 25, 26 and 28 appear to arise from the same misunderstanding. Applicant refers the Office to the section discussing those claims below.

Objection to claims 14 and 29

Applicant amends both claims to insert the conjunctive “and” before the last step recited in each claim.

In both claims, the phrase “a competing-algorithm” in “a competing-algorithm performance of said competing algorithm” is a compound adjective that modifies the noun “performance.” Its presence makes clear that the performance referred to is that of the competing algorithm, and not that of the incumbent algorithm.

Applicant requests withdrawal of the objection insofar as it requires removal of the compound adjective “competing-algorithm.”

Section 112 rejection: “preferable”

The Office appears to consider that the term “preferable” in claims 2 and 17 renders those claims vague and indefinite because it is not clear what standard is used to determine that one algorithm is preferable to another.

In response, Applicant amends these claims to recite providing data indicative of a *performance difference* between the competing algorithm and the incumbent algorithm.

Section 112 rejection: “selected interval”

The Office appears to consider the term “a selected interval” in claims 3, 4, 5, 7, 18, 19, 20, and 22 to be unclear because it is not apparent whether the interval is a time interval or some other type of interval.

At the time of drafting, Applicant briefly considered and rejected the idea of including an adjective such as “time” or “temporal” to modify “interval.” Applicant concluded that the

adjective would be superfluous verbiage because, in the context of the specification, no one could possibly construe the interval to be anything but a time interval.

Applicant also draws attention to the use of “during” in all the foregoing claims except claims 3 and 18. The preposition “during” is one that frequently connotes a time interval rather than any other kind of interval.

Claims 3 and 18 are amended to recite “during” instead of “over,” thereby rendering them consistent with the remaining claims.

Applicant submits that it is quite clear, both from the context of the specification and from the use of the preposition “during,” that the recited “interval” could only be a time interval.

Section 112 rejection: “cost”

The Office appears to consider the language “cost associated with replacing said incumbent algorithm” in claims 9, 13, 24, and 28 to be vague and indefinite because it is unclear whether “cost” relates to finance or performance.

In response, Applicant draws attention to the distinction between cost and the specific ways of measuring cost. The term “cost” is intended to include any kind of cost, without being limited to specific ways of measuring cost.

Section 112 rejection: “simulating performance”

The Office appears to consider “simulating a performance of said competing algorithm” in claims 1, 10, 11, 13, 16, 21, 25, 26 and 28 to be indefinite because it is unclear whether “performance” means “execution” or “efficiency.”

Applicant amends claim 1’s third paragraph to recite a “corresponding performance of said competing algorithm.” This amendment makes it clear that the “performance” recited in the second paragraph can be none other than a measure of how well a competing algorithm performed. Note that in the second paragraph, it is already clear that “performance” cannot

possibly mean “execution” because that paragraph already refers to the competing algorithm as “executing in place of said incumbent algorithm.”

The foregoing remarks apply to claim 16, which recites limitations similar to those of claim 1.

Claims 10 and 25 are believed to be require no amendment because the use of the subjunctive in “simulating a performance of said competing algorithm *were it to execute...* in place of the incumbent algorithm” makes the distinction between the “performance” as a measure of an outcome and “execution” as what results in that measure of performance is already clear.

Claims 11 and 26 are amended to recite “said performance of said competing algorithm” thereby causing that performance to clearly refer to the performance referred to in paragraph 2 of claims 10 and 25, from which claims 11 and 26 depend. Since the meaning of “simulating a performance” is clear in claims 10 and 25, this amendment to claims 11 and 26 results in those claims inheriting that clarity.

Applicant amends claim 13 to recite “said simulated performance,” thereby causing that claim to refer explicitly to the performance simulated in the second paragraph of claim 11 from which it depends. A similar amendment is made to claim 28.

Claim 21 is amended so that the second step explicitly refers to the second step of claim 16 from which it depends. Since that step is clear in claim 16, claim 21 inherits that clarity.

The Office appears troubled by the notion of “simulating a number.” In fact, it is not a number that is being simulated, but a performance that results from having executed an algorithm. A number, such as the “75%” referred to in the office action, is merely an outcome of the simulation.

By way of analogy, one might simulate the performance a heating system by creating a mathematical model and assuming a particular weather pattern. The resulting fuel consumption, which is a number, would represent the simulated performance of the heating system.

Applicant takes the opportunity to draw a distinction between “simulation of a performance” and “evaluation of a performance” as those terms are used in the claims.

When *simulating* the performance of a heating system, one need not have an actual system. The simulation process is carried out by creating a model. In contrast, *evaluating* the performance of a system would mean observing the actual outcome of using the system. In evaluating the claims, Applicant requests that the Office carefully consider this distinction.

Section 112 rejection: “data indicative of a comparison”

Claims 14 and 29 appear to be rejected because the absence of criteria for “a comparison between said incumbent algorithm and said competing algorithm” allegedly renders this step unclear.

Applicant submits that there is no need to specify particular criteria for comparison since to do so would unnecessarily limit the claim to those specific criteria.

The Examiner has correctly interpreted the limitation as being “data indicative of performance.” Accordingly, Applicant amends the claims in a manner consistent with the Examiner’s interpretation.

Section 112 rejection: “a performance selected from said actual performance”

In claim 13, the Office appears troubled by the language “ a performance selected from said actual performance.” Applicant points out that this claim recites a Markush group having two elements: actual performance and simulated performance. The selected performance can thus be either the simulated performance or the actual performance.

To assist the Office, the claim is amended to draw attention to the Markush group. The claim is also reformatted to draw attention to its grammatical structure.

Section 112: lack of antecedent basis

Applicant amends claims 14, 5, and 20 to introduce antecedent basis as needed.

Section 101 rejection

In response to the section 101 rejection of claim 15, Applicant amends claim 15 to more clearly recite statutory subject matter.

Section 102 rejection of claim 1

As best understood, the Office considers claim 1's step of

evaluating an incumbent-algorithm score indicative of a performance of an incumbent algorithm

to correspond to *Cooper*'s disclosure,¹ in FIG. 3, of creating a baseline network simulation. It is apparently the Office's position that a "network simulation" corresponds to "an incumbent algorithm."

As a threshold matter, Applicant submits that a network and an algorithm are two completely different things. A *network* is a physical structure for transmission of data. It includes routers and cables. An *algorithm* is a sequence of steps for manipulating data. It is unclear on what basis a network and an algorithm can possibly be considered even close to the same thing.

Assuming, for the sake of argument, that a "network" is an "algorithm," *Cooper* nevertheless fails to teach the claimed subject matter.

Cooper discloses a system for simulating the effectiveness of a network under a variety of conditions. In *Cooper*, a user generates a "baseline network simulation" by first defining the network to be simulated. This includes specifying the network architecture and any user preferences. Then, the user defines a model of the traffic that the network is expected to bear. The user then instructs simulation software to simulate the effectiveness of the network under those user-defined conditions.

In particular, *Cooper* states that

¹ *Cooper, et al.*, U.S. Patent No. 5,809,282

"The windows of graphical user interface 240 are the "canvas" on which the user creates simulations of a network based on (1) a specific network configuration and (2) defined traffic conditions.

The configuration of a network includes information on the location of sites (e.g., nodes) in the network and links connecting the sites (i.e., connectivity), as well as the hardware used at each site and for each link. The traffic conditions of the network, also called a "scenario," include information on the type and size of transmissions in the network along with the time when the transmissions occur and the source and destination sites for each transmission. (From the network configuration and traffic conditions, network simulator 210 generates a simulation, including the performance and cost of operating the network under the defined traffic conditions.)

The user also inputs, via graphical user interface 240, preference data related to qualitative concerns that the user may have with regard to operation of the network. Preference data includes whether a network operator would like to down-size its network; decrease costs of using, operating, or maintaining the network; update network equipment; minimize transmission delays; minimize the number of messages lost during transmission; maximize network utilization at all times; minimize network changes but satisfy other preferences and/or performance requirements; and make sure that state-of-the-art equipment is employed in the network.

Network simulator 210 is initiated by the user using graphical user interface 240 to generate a network simulation from the specific network configuration and defined traffic conditions. A network simulation, called a base-line simulation, includes data related to effectiveness (performance and cost) of operation of the network under the defined traffic conditions."²

Applicant draws attention to the fact that *Cooper* teaches *simulating* a performance of a network. *Simulating* a performance of a network is different from actually *evaluating* a performance of a network. A network whose performance is being *simulated* does not have to actually exist.

In fact, in many cases, the whole purpose of simulating a network is to decide whether that network is actually worth building. There would be no point in simulating a network that already existed. One could simply go out and evaluate its actual performance.

Claim 1's first step recites "*evaluating* an incumbent-algorithm score;" it does *not* recite "*simulating* an incumbent-algorithm score." The verb "simulating" is saved for the second step, which recites "simulating performance of a competing algorithm executing in place of said incumbent algorithm." This choice of words was no accident.

² *Cooper*, col. 5, line 49-col. 6, line 16.

In Applicant's system, the performance of an incumbent-algorithm is *not* being simulated. It does *not* execute on manufactured data within the dry confines of some research laboratory. In Applicant's system, the incumbent-algorithm score is evaluated while the incumbent algorithm is actually executing in the real world, with all its unpredictability, using real data provided by live users.

In contrast, the performance of the competing algorithms *is* being simulated. These algorithms are not actually executing on the real data. To do so would wastefully consume bandwidth, which would in turn adversely affect the performance of the incumbent algorithm. Instead, the performances of the competing algorithms are obtained by determining what *would have* resulted had those algorithms been executing on the same data provided to the incumbent server.

Applicant's invention thus provides a way to monitor the effectiveness of several algorithms, one of which is the incumbent algorithm, as they are subjected to the same demands made by actual users. It provides a way to answer the question "what if a different algorithm had been used on this exact same data stream." The invention thus enables an administrator of a data storage system to adaptively react in real time to changes in how users use the data storage system.

Cooper teaches nothing at all like this. There is no suggestion in *Cooper* of evaluating the performance of one "incumbent" network architecture as it carries live network traffic, with all its inherent unpredictability, and then simulating the performance of other "competing" network architectures on the identical traffic.

What *Cooper* teaches is a way to *simulate* several different network architectures using the same *pre-defined models* of network traffic. One would hope that these pre-defined models of traffic would bear some relationship to actual traffic. However, in the final analysis, the pre-defined traffic models are only models; they are not real.

That *Cooper* only discloses simulating, and not evaluating, network performance is apparent from FIG. 3, to which the Office has already drawn attention. Step **320** is the step of creating a baseline network *simulation*, i.e. simulating an “incumbent” network architecture. Step **350** is that of creating new network simulations for different network modifications, i.e. simulating “competing” network architectures. There is nothing in FIG. 3 that requires that an actual network be available. Apparently, what *Cooper* discloses can be carried out without being connected to any network at all.

Accordingly, the section 102 rejection of claim 1 is deficient and should be withdrawn. The remaining independent claims recite limitations similar to those of claim 1. Accordingly, the section 102 rejections of those claims should also be withdrawn.

The remaining claims are dependent claims. Those claims are therefore allowable for at least the same reasons as the independent claims from which they depend.

Now pending in this application are claims 1-29, of which claims 1, 10, 14, 15, 16, and 29 are independent. No fees are believed to be due in connection with the filing of this response. However, to the extent fees are due, or if a refund is forthcoming, please adjust our deposit account 06-1050.

Respectfully submitted,

Date: January 25, 2005



Faustino A. Lichauco
Reg. No. 41,942

Fish & Richardson P.C.
225 Franklin Street
Boston, MA 02110-2804
Telephone: (617) 542-5070
Facsimile: (617) 542-8906